

Attorney Docket No. 16869P-105400US
Client Ref. No. 340301085US01

PATENT APPLICATION

DISK ARRAY OPTIMIZING THE DRIVE OPERATION TIME

Inventors:

Ikuya YAGISAWA
Citizenship: Japan

Naoto MATSUNAMI
Citizenship: Japan

Akihiro MANNEN
Citizenship: Japan

Kenichi TAKAMOTO
Citizenship: Japan

Assignee:

Hitachi, Ltd.
6, Kanda Surugadai 4-chome
Chiyoda-ku, Tokyo, Japan
Incorporation: Japan

Entity:

Large

TOWNSEND AND TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
(415) 576-0200

TITLE OF THE INVENTION

DISK ARRAY OPTIMIZING THE DRIVE OPERATION TIME

BACKGROUND OF THE INVENTION

The present invention relates to an external storage device system. More specifically, the present invention relates to a technology for prolonging an operation period of a disk device (hereafter also referred to simply as a disk) and decreasing power consumption of a storage device system (hereafter referred to as a disk array). Here, the disk device's operation period signifies a period from the time to start using the disk device to the time when the disk device becomes unusable.

A disk array is a type of storage device systems connected to a computer. The disk array is also referred to as a RAID (Redundant Arrays of Inexpensive Disks) and constitutes a storage device system comprising a plurality of disk devices arranged in an array and a control section to control them. The disk array concurrently operates disk devices to accelerate read requests (requests to read data) and write requests (requests to write data) and to provide data with redundancy. Disk arrays are categorized into five levels depending on types of redundant data to be added and disk array configurations. This is described in non-patent document 1 (Daved A. Patterson, Garth Gibson, and Randy H. Katz, "A Case for Redundant Arrays of Inexpensive Disks (RAID)", Computer Science Division Department of Electrical

Engineering and Computer Sciences, University of California Berkeley, 1988).

SUMMARY OF THE INVENTION

Some disk devices used for a disk array may shorten a total active time when they are turned on 24 hours a day in comparison to a case where the disk devices are turned on only when needed. The total active time signifies the accumulated time during which the disk device is turned on. Such problem caused by deterioration and wear of parts constituting the disk devices. The disk device's total active time shortens when the disk device is turned on 24 hours a day and is operated continuously 24 hours a day. As a result, the disk device's operation period shortens. Such disk device is not suited for 24-hour continuous operations. The disk array capacity increases year after year. There is a trend of increasing the number of disk devices to be mounted. The power consumption for disk arrays tends to increase.

The present invention discloses a technology to prolong operation times of disk devices constituting a disk array. The present invention also discloses a technology to decrease the disk array's power consumption.

A computer accessing a disk array has a disk control instruction program which allows the disk array to turn on or off specific disk devices in it. The disk array has a disk power supply control instruction reception program and

a disk power supply control circuit. The disk power supply control instruction reception program receives an instruction to turn on or off disk devices from the computer. The disk power supply control circuit turns on or off a disk device corresponding to an LU specified by the computer.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows an example of a system configuration according to a first embodiment;

FIG. 2 shows an example of a computer configuration according to the first embodiment;

FIG. 3 shows an example of a backup server configuration according to the first embodiment;

FIG. 4 shows an example of a disk management table according to the first embodiment;

FIG. 5 shows an example of an update position management table according to the first embodiment;

FIG. 6 shows an example of obtaining or deleting a snapshot according to the first embodiment;

FIG. 7 shows a process example of a mirror resynchronization program according to the first embodiment;

FIG. 8 shows an example of a system configuration according to a second embodiment;

FIG. 9 shows an example of a computer configuration according to the second embodiment;

FIG. 10 shows an example of an address map according to the second embodiment; and

FIG. 11 shows a process example of a disk power supply control instruction program according to the second embodiment.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Embodiments of the present invention will be described in further detail with reference to the accompanying drawings.

First embodiment

The first embodiment will now be described.

(1) Description of the system configuration

FIG. 1 shows an example of the system configuration according to the first embodiment. In FIG. 1, a computer 100 is connected to a disk array 200 via a Fibre Channel interface (hereafter referred to as an FC I/F) 290 of the disk array 200. A backup server 400 is connected to the disk array 200 via an FC I/F 300 of the disk array 200. A tape apparatus 500 is connected to the backup server 400. A management terminal 600 is connected to the disk array 200 via a management I/F 310 of the disk array 200.

The disk array 200 comprises a CPU 210, memory 220, a cache 240, and a disk controller 250. The CPU 210 controls the disk array 200. The cache 240 stores user data of the disk array 200. The disk controller 250 controls a plurality of disk devices.

The disk controller 250 is connected to a plurality of FC disks 271 through 273 (hereafter referred to as an FC disk group) and a plurality of ATA disks 274 through 276

(hereafter referred to as an ATA disk group). The FC disk group and the ATA disk group each use a redundant configuration called RAID. The computer 100 accesses storage areas in the FC disks 271 through 273 as a SCSI logical unit (LU). Likewise, the computer 100 also accesses storage areas in the ATA disks 274 through 276 as a SCSI LU. An LU 261 represents the LU in the FC disks 271 through 273. An LU 262 represents the LU in the ATA disks 274 through 276. FIG. 1 illustrates the LU 261 and the LU 262 each as a storage area covering a plurality of disks. Further, the LU may comprise a storage area on one disk.

According to the embodiment, the snapshot management program 226 in the disk array 200 manages the LU 261 and the LU 262 duplicatively. When data is written to the LU 261 according to the embodiment, the snapshot management program 226 writes the same data to the LU 262 to duplicate the data. In this manner, storing the same data in the LU 261 and the LU 262 is referred to as mirroring data. The LU 261 contains original data and is referred to as a mirror origin LU. The LU 262 contains a copy of the original data and is referred to as a mirror destination LU. The LU 262 is used to obtain a snapshot. The snapshot is used to back up data on the tape apparatus.

A disk power supply control circuit 280 is connected to the ATA disks 274 through 276. The disk power supply control circuit 280 turns on or off each ATA disk independently of the other ATA disks. The disk power supply control circuit

280 is provided with registers corresponding to the ATA disks 274 through 276. Turning on or off the register turns on or off the corresponding ATA disk.

The following describes programs and management tables in the disk array 200.

The memory 220 of the disk array 200 contains a RAID control program 221, a disk management program 222, a disk power supply control instruction reception program 223, a disk management table 224, a disk power supply control program 225, and a snapshot management program 226. The RAID control program 221 controls the disk array 200. The disk management program 222 manages the FC disks 271 through 273 and the ATA disks 274 through 276. The disk management table 224 records operation parameters and operation states of the FC disks 271 through 273 and the ATA disks 274 through 276. The disk power supply control instruction reception program 223 receives an instruction to turn on or off disk devices from the computer 100. This instruction is hereafter referred to as a disk power supply control instruction. The disk power supply control program 225 turns on or off the ATA disks 274 through 276 based on instructions received from the computer 100. The snapshot management program 226 controls snapshots. According to the embodiment, a snapshot instruction program 127 to be described is contained in the computer 100 and issues instructions to delete or obtain snapshots. In addition, the snapshot instruction program 127 issues ModeSelect commands concerning disk power supply control instructions

for the ATA disks 274 through 276. Accordingly, the snapshot management program 226 accepts not only instructions to delete or obtain snapshots, but also disk power supply control instructions for the ATA disks 274 through 276.

The snapshot management program 226 also contains a disk access sub-program 227 and an LU mirror sub-program 228. The disk access sub-program 227 instructs the disk controller 250 to access disks in accordance with requests from the computer 100. The LU mirror sub-program 228 applies an update to one LU and another predetermined LU and writes the same user data to two LUs. The LU mirror sub-program 228 duplicates data written to the LU 261 for the LU 262.

The snapshot management program 226 further contains a non-mirror update monitoring sub-program 230, a non-mirror update position management sub-program 231, and a mirror resynchronization sub-program 229. When the data duplication is inactivated between the LU 261 and the LU 262 (non-mirror state), the non-mirror update monitoring sub-program 230 detects an update to the mirror origin LU. The non-mirror update position management sub-program 231 records the update position in an update position management table 232 to be described later. When the mirror resynchronization is performed to restart data duplication between the LU 261 and the LU 262, the mirror resynchronization sub-program 229 copies the updated part of the mirror origin LU to the mirror destination LU.

The update position management table 232 is used to

manage the contents of data stored in the mirror origin LU and the mirror destination LU. When data in the mirror origin LU is updated in the non-mirror mode, the update position management table 232 records the update position in the mirror origin LU.

For example, the update position management table 232 is a bit map as shown in FIG. 5 and lists LBA set numbers in the mirror origin LU and update bit states corresponding to the LBA set numbers. Each LBA set comprises one, or the same number of LBAs (Logical Block Addresses) that are used as a unit to divide the entire area in the LU from the beginning. The LBA sets are assigned LBA set numbers, i.e., sequence numbers, from the beginning.

The update bit indicates whether or not the LBA set corresponding to the update bit is updated in the non-mirror mode. The update bit is set to 1 when the LBA set is updated, or to 0 otherwise. The update bit is initially set to 0. For example, the update position management table 232 in FIG. 5 shows that only the area corresponding to LBA set number 1 is updated in the non-mirror mode.

Referring now to FIG. 4, an example of the disk management table 224 will now be described. The disk management table 224 comprises the following columns. The "diskNo." column shows identification numbers of disk devices in the disk array 200. The "disk type" column shows the disk typ , i.e., whether the disk device is an FC disk or an ATA disk. The "array configuration" column shows to which RAID

group the disk device belongs. The "LU No." column shows to which LU the disk device belongs. The "snapshot pair" column shows whether data stored in the disk is original or copy. The "snapshot status" column shows whether or not data to be stored in an FC disk is mirrored to an ATA disk. The "disk startup status" column shows whether the disk device is turned on or off. The "accumulated time" column shows the accumulated time in which the disk device is turned on. The "life setting" column shows the limit of the accumulated time in which the disk device is turned on.

Returning now to FIG. 1, the management terminal 600 comprises an input section 610 and an output section 620. The input section 610 accepts settings for the FC disks 271 through 273 and the ATA disks 274 through 276 from a user. The output section 620 displays information about the FC disks 271 through 273 and the ATA disks 274 through 276 to the user.

The input section 610 inputs parameters to be assigned to the disk management table 224. The output section 620 outputs information recorded in the disk management table 224. The input section 610 represents a mouse, a keyboard, and the like. The output section 620 represents a display, an liquid crystal panel, a buzzer, and the like.

Referring now to FIG. 2, an example of the computer 100 will now be described. The computer 100 comprises a CPU 110, an FC I/F 140, a communication I/F 160, and memory 120. The CPU 110 controls the computer 100. The I/F 140 connects the computer 100 to the disk array 200. The communication

I/F 160 provides communication between the computer 100 and the backup server 400. The memory 120 stores programs.

The memory 120 stores a database program 126, a snapshot instruction program 127, and a disk power supply control instruction program 128. The snapshot instruction program 127 allows the disk array 200 to obtain snapshots. The disk power supply control instruction program 128 instructs to turn on or off disk devices in the disk array 200. The CPU 110 to control the computer 100 executes these programs. The snapshot instruction program 127 generates and issues SCSI ModeSelect commands. The ModeSelect commands include an instruction to turn on or off the ATA disks 274 through 276 from the disk power supply control instruction program 128 to the disk array 200. The ModeSelect commands also include an instruction to delete or obtain snapshots from the snapshot instruction program 127 to the disk array 200. The database program 126 accesses the LU 261, i.e., the mirror origin LU while the database is referenced or updated. The database program 126 also controls data update. The database program 126 stops accessing the disk array 200 at a check point that indicates a meaningful separation as database information. When data in the LU 261 becomes meaningful as database information, the database program 126 changes the disk array to a backup mode to back up that data. The backup mode ensures the consistency of data in the LU 261.

Referring now to FIG. 3, an example of the backup server 400 will now be described. The backup server 400 comprises

a CPU 410, an FC I/F 440, a communication I/F 460, a SCSI I/F 450, and memory 420. The CPU 410 controls the backup server 400. The FC I/F 440 connects the backup server 400 with the disk array 200. The communication interface 460 provides communication between the backup server 400 and the computer 100. The SCSI I/F 450 connects the backup server 400 with the tape apparatus. The memory 420 stores programs. The memory 420 stores a backup program 426, an I/O control program 427, and a tape control program 428. The backup program 426 backs up data in the disk array 200 onto the tape apparatus 500. The I/O control program 427 transfers data between the disk array 200 and the tape apparatus 500. The tape control program 428 control the tape apparatus 500. The CPU 410 to control the backup server 400 executes these programs.

There has been described the system configuration of the embodiment.

(2) Obtaining and deleting snapshots

With reference to the flowchart in FIG. 6, the following describes an example of processes performed by the snapshot instruction program 127 and the snapshot management program 226 when snapshots are obtained or deleted. The embodiment assumes that the disk array 200 obtains a snapshot of the LU 261 in synchronization with the timing when the backup server 400 backs up data. It also assumes to turn on the ATA disks 274 through 276 for the LU 262 only during a period in which a snapshot is obtained and data is backed up from

mirror destination LU.

Since the LU 262 stores data of the previously obtained snapshot, the snapshot instruction program 127 running on the computer 100 needs to nullify that snapshot and obtain a new snapshot. Hereafter, this is referred to as deletion of the snapshot. For this purpose, the snapshot instruction program 127 issues a ModeSelect command to the disk array 200 to delete the snapshot (step 2001).

The snapshot management program 226 in the disk array 200 receives the ModeSelect command (step 3001). Based on the disk power supply control instruction included in the ModeSelect command, the snapshot management program 226 turns on the ATA disks 274 through 276 (step 3002). The disk power supply control instruction turns on the disk devices corresponding to the LU 262. According to the instruction, the disk array 200 locates ATA disks constituting the LU 262 from the disk management table 224 in FIG. 4 and turns on the ATA disks 274 through 276 in this embodiment. The LU mirror sub-program 228 is activated to restart duplication of the LU 261 and the LU 262 (step 3003). In this embodiment, the activation means starting or initiating a process or a program. The inactivation means terminating or stopping a process or a program. At step 3003, the LU 262 reflects the update to the LU 261.

The snapshot management program 226 inactivates the non-mirror update monitoring sub-program 230 and the non-mirror update position management sub-program 231 to stop

recording updates to the LU 261 (step 3004). Thereafter, the update bit in the update position management table 232 is not changed by the non-mirror update position management sub-program 231.

The snapshot management program 226 then activates the mirror resynchronization sub-program 229. The mirror resynchronization sub-program 229 references the update position management table 232. When the LU 261 contains data that differs from data stored in the LU 262, the mirror resynchronization sub-program 229 copies that data from the LU 261 to the LU 262 to resume the mirror state between the LU 261 and LU 262 (step 3005). This is also referred to as resynchronization.

Upon completion of the resynchronization, the snapshot management program 226 inactivates the mirror resynchronization sub-program 229 (step 3006). The snapshot management program 226 sends the termination status of the ModeSelect command to the snapshot instruction program 127 on the computer 100 (step 3007). The snapshot instruction program 127 receives the termination status of the ModeSelect command to terminate the operation (step 2002).

The snapshot instruction program 127 on the computer 100 then issues a ModeSelect command to obtain a snapshot on the disk array 200 (step 2003).

The snapshot management program 226 on the disk array 200 receives the ModeSelect (step 3008). The snapshot management program 226 activates the non-mirror update

monitoring sub-program 230 and the non-mirror update position management sub-program 231 to allow the update position management table 232 to record the position information about the data update applied to the LU 261 (step 3009). That is to say, when the LU 261 is updated, the update position management table 232 indicates the update bit set to 1 for the LBA set including the updated LBA to record the update.

The snapshot management program 226 then inactivates the LU mirror sub-program 228 to stop duplication of the LU 261 and the LU 262 (step 3010). This prevents the LU 262 as the mirror destination LU from reflecting the update to the LU 261 as the mirror origin LU.

The snapshot management program 226 then sends the termination status of the ModeSelect command to the snapshot instruction program 127 on the computer 100 (step 3011).

The snapshot instruction program 127 on the computer 100 receives the termination status of the ModeSelect command from the snapshot management program 226 (step 2004). The snapshot instruction program 127 supplies the backup server 400 with an instruction to obtain a backup via the communication I/F 160. The backup server 400 receives the termination report (step 2005). The snapshot instruction program 127 then issues a ModeSelect command to turn off the ATA disks 274 through 276 to the snapshot management program 226 (step 2006). The procedure to turn off the ATA disks 274 through 276 is the same as that at step 3002.

When receiving the ModeSelect command, the snapshot

management program 226 in the disk array 200 turns off the ATA disks 274 through 276 according to the disk power supply control instruction contained in the ModeSelect command (step 3012).

There have been described the operations of the snapshot instruction program 127 and the snapshot management program 226 when snapshots are obtained or deleted.

(3) Write operation

The following describes a write operation of the RAID control program 221 in the disk array 200. When the computer 100 writes data to the LU 261 to update the contents stored in the LU 261, the snapshot management program 226 receives a WRITE command and data issued to the LU 261.

Let us assume that the LU mirror sub-program 231 is active and that the non-mirror update monitoring sub-program 230 and the non-mirror update position management sub-program 231 are inactive. When receiving write data from the computer 100, the snapshot management program 226 writes the write data to both the LU 261 as the mirror origin LU and the LU 262 as the mirror destination LU.

Let us assume that the LU mirror sub-program 231 is inactive and the non-mirror update monitoring sub-program 230 and the non-mirror update position management sub-program are active. In this case, the snapshot management program 226 writes data to the LU 261 as the mirror origin LU. While the update position management table 232 records update positions, the snapshot management program 226 allows this

table to set the update bit to 1 for the LBA set including the updated LBA.

Finally, the status is sent to the computer 100.

(4) Mirror resynchronization

The following describes an example of the mirror resynchronization to copy data from the LU 261 to the LU 262 at step 3003 in FIG. 6. FIG. 7 exemplifies an operation of the mirror resynchronization sub-program 229 during mirror resynchronization. The mirror resynchronization sub-program 229 checks whether or not the update position management table 232 contains the update bit set to 1 to record the update (step 1001). When there is no update bit set to 1 to record the update, the mirror resynchronization is complete. The mirror resynchronization sub-program 229 terminates the process (step 1002).

When there is an update record, the mirror resynchronization sub-program 229 references the update position management table 232. Based on the LBA set number corresponding to the update bit set to 1, the mirror resynchronization sub-program 229 calculates a read position on the LU 261 and a write position on the LU 262 (step 1003).

The mirror resynchronization sub-program 229 prevents data from being updated to the corresponding read/write position (step 1004). Based on the calculated read/write position, the mirror resynchronization sub-program 229 copies the data stored at the read position in the LU 261 as the mirror origin LU to the write position in the LU 262

as the mirror destination LU (step 1005). It should be noted that preventing the update is equivalent to inhibiting writing. The update is prevented to ensure the consistency of data between the LU 261 and the LU 262 after the mirror resynchronization.

The mirror resynchronization sub-program 229 then releases the prevention against updating data at the corresponding read/write position (step 1006). The mirror resynchronization sub-program 229 sets 0 to the update bit in the update position management table 232 correspondingly to the mirrored position due to the copy at step 1005 (step 1007), and then returns to step 1001.

There has been described the mirror resynchronization.

(5) Disk operation time management

Using the input section 610 of the management terminal 600, a user can set the life of each of the FC disks 271 through 273 and the ATA disks 274 through 276 in the disk management table 224. The life signifies a product cycle estimated from the disk design specifications of the disk device, or a warranty period, or an accumulated operation time over which the disk is highly possibly subject to errors. It is assumed to find an accumulated time by totaling the time during which the disk device is turned on. When the accumulated time exceeds the life, there is a high possibility of causing an error.

The disk management program 222 records the disk device's accumulated time in the disk management table 224.

The disk management program 222 counts the time during which the FC disks 271 through 273 and the ATA disks 274 through 276 are turned on. The disk management program 222 records the counted time as the accumulated time in the disk management table 224. The disk management program 222 compares the life of each of the FC disks 271 through 273 and the ATA disks 274 through 276 with the accumulated time. When the accumulated time exceeds the life, the disk management program 222 displays an alarm on the output section 620 of the management terminal 600.

According to the first embodiment as mentioned above, the snapshot management program obtains a snapshot of the LU 262 in synchronization with the timing when the backup server obtains a backup. It is possible to turn on the ATA disk device storing the snapshot data only when the snapshot management program obtains a snapshot and backs it up on the tape. In other words, the disk device can be turned off while no backup is created. This can extend the operation period of disk devices in the disk array 200 and decrease the power consumption of the disk array 200.

The disk array 200 receives the ModeSelect command issued from the computer 100. At this time, the disk controller 250 controls turning on or off the disk device. While the disk device is turned off, the computer 200 does not access the disk device. This can prevent a time-out condition of access requests from the computer 200.

Second embodiment

The second embodiment will now be described.

(1) Description of the system configuration

FIG. 8 shows an example of the system configuration according to the second embodiment. The following describes only differences from the first embodiment. Unlike the example in FIG. 1, the system in FIG. 8 uses entirely ATA disks 281 through 289 connected to the disk controller 250. The computer 100 accesses a storage area in the ATA disks 281 through 283 as a SCSI LU. Likewise, the computer 100 accesses a storage area in the ATA disks 284 through 286 as another SCSI LU. The computer 100 accesses a storage area in the ATA disks 287 through 289 as yet another SCSI LU. The LU in the ATA disks 281 through 283 is defined as an LU 263. The LU in the ATA disks 284 through 286 is defined as an LU 264. The LU in the ATA disks 287 through 289 is defined as an LU 265. FIG. 8 shows that each of the LU 263, the LU 264, and the LU 265 constitutes a storage area extending to a plurality of disks. Each LU may comprise a storage area in one disk. According to an instruction from the disk power supply control program 225, the disk power supply control circuit 280 turns on or off the ATA disks 281 through 289 independently of the other ATA disks.

The memory 220 stores an LU protection program 234 that protects the LUs against reading or writing instead of the snapshot management program 226.

Unlike the example in FIG. 4, the disk management table 224 according to the second embodiment does not record the

snapshot pair or the snapshot status.

FIG. 9 shows an example of the computer 100 according to the second embodiment. Unlike the example in FIG. 2, the computer 100 contains an E-mail application program 130, an LU protection instruction program 131, and a disk power supply control instruction information storage area 132 instead of the database program 126 and the snapshot instruction program 127. The LU protection instruction program 131 instructs protection against reading or writing to the LUs. The E-mail application stores electronic mail archives in the memory 120, sequentially writes data, and stores data at consecutive addresses. This embodiment will be described using the E-mail application as an example. The E-mail application specifies a sequence of LUs to be written. According to the embodiment, the E-mail application accesses the LU 263, LU 264, and LU 265 in this order at scheduled times to write data. The embodiment is not limited to the E-mail application.

The disk power supply control instruction information storage area 132 stores an address map as shown in FIG. 10. The address map lists LU numbers, LBA numbers, and disk startup status of disk devices constituting LUs. The disk startup status provides information about power-on/off states of the disk devices.

The computer further contains a management I/F 150 that is connected to the management terminal 600.

Th r has been described the system configuration

according to the second embodiment.

(2) Operation of the disk power supply control instruction program

The E-mail application 130 issues a request to write data to the disk array 200 for sequential access. Based on this request, the disk power supply control instruction program 128 in the computer 100 references the address map in FIG. 10. The disk power supply control instruction program 128 issues an instruction to turn on disk devices belonging to the LU corresponding to the address for writing. The disk power supply control instruction program 128 issues an instruction to turn off disk devices belonging to the LU corresponding to the address for which the writing is complete. When reading data, a user uses the input section 610 of the management terminal 600 to issue an instruction to turn on disk devices belonging to the LU that stores data to be read.

FIG. 11 is a flowchart showing an operation example of the disk power supply control instruction program 128 in the E-mail application. Let us assume that the E-mail application 130 makes preparations for writing to the LU 263. Based on a request from the E-mail application 130 to write to the disk array 200, the disk power supply control instruction program 128 instructs the disk array 200 to turn on the disks constituting the LU 263 (step 5001). The E-mail application 130 starts writing to the LU 263. The disk power supply control instruction program 128 then determines whether or not the remaining time is shorter than or equal

to the specified time until the scheduled time of terminating the access to the LU 263 (step 5002). When the remaining time is longer than the specified time until the scheduled time of terminating the access, the disk power supply control instruction program 128 repeats the determination at step 5002. When the remaining time becomes shorter than the specified time, the disk power supply control instruction program 128 instructs the disk array 200 to turn on the LU 264 (step 5003). The disk power supply control instruction program 128 then determines whether or not the E-mail application 130 starts accessing the LU 264 (step 5004). When the E-mail application 130 does not start accessing the LU 264, the disk power supply control instruction program 128 repeats the determination at step 5004. When the E-mail application 130 starts accessing the LU 264, the disk power supply control instruction program 128 instructs the disk array 200 to turn off the LU 263 that terminated the access from the E-mail application 130 (step 5005). The disk power supply control instruction program 128 then determines whether or not the remaining time is shorter than or equal to the specified time until the scheduled time of terminating the access to the LU 264 (step 5006). When the remaining time is longer than the specified time until the scheduled time of terminating the access, the disk power supply control instruction program 128 repeats the determination at step 5006. When the remaining time becomes shorter than the specified time, the disk power supply control instruction

program 128 instructs the disk array 200 to turn on the LU 265 (step 5007). The disk power supply control instruction program 128 then determines whether or not the E-mail application 130 starts accessing the LU 265 (step 5008). When the E-mail application 130 does not start accessing the LU 265, the disk power supply control instruction program 128 repeats the determination at step 5008. When the E-mail application 130 starts accessing the LU 265, the disk power supply control instruction program 128 instructs the disk array 200 to turn off the LU 264 that terminated the access from the E-mail application 130 (step 5009). Finally, the disk power supply control instruction program 128 determines whether or not the access to the LU 265 terminates (step 5010). When the access does not terminate, the disk power supply control instruction program 128 repeats the determination at step 5010. When the access from the E-mail application 130 terminates, the disk power supply control instruction program 128 instructs the disk array 200 to turn off the LU 265 (step 5011). There has been described the operation of the disk power supply control instruction program 128. Like the first embodiment, the disk array 200 receives a disk power supply control instruction from the computer 100. The disk array 200 references the disk management table 224 to locate the disk device corresponding to the LU to be turned on or off and turns on or off the disk device.

It may be preferable to instruct the disk array 200 to turn on or off disk devices corresponding to the LU based

on a user input. For example, the user may use the input section 610 of the management terminal 600 to issue an instruction to the LU protection instruction program 131 so as to protect an LU against reading or writing. In this case, the disk power supply control instruction program 128 instructs the disk power supply control program 225 of the disk array 200 to turn off disk devices belonging to the LU that is instructed to be protected. The user may use the input section 610 of the management terminal 600 to issue an instruction to the LU protection instruction program 131 so as to release the protection of an LU against reading or writing. In this case, the disk power supply control instruction program 128 instructs the disk power supply control program 225 of the disk array 200 to turn on disk devices belonging to the LU that is instructed to be unprotected.

The embodiment has been described using the E-mail application as an example. Consequently, the flowchart in FIG. 11 shows the operation of specifying the sequence of LUs for writing data and sequentially turning on or off the disk devices belonging to the LU at specified times. However, the embodiment is not limited to the E-mail application and is capable of turning on or off disk devices constituting the LU to be accessed for reading or writing.

The embodiment can operate only disk devices reading or writing data and stop the other disk devices not reading or writing data, making it possible to extend the disk array's

operation period and decrease its power consumption.

The disk controller 250 controls turning on or off the disk devices in synchronization with the timing when the disk array 200 receives an instruction from the computer 100 to turn on or off the disk devices. While the disk device is turned off, the computer 200 does not access the disk device. This can prevent a time-out condition of access requests from the computer 200.

It is possible to extend the operation period and the total active time of the disk devices constituting the disk array and decrease the disk array's power consumption.